arXiv:1607.03479v2 [cs.SY] 14 Jul 2016

SAT-based Distributed Reactive Control Protocol Synthesis for Boolean Networks

Yunus Emre Sahin, Necmiye Ozay

***Abstract*— This paper considers the synthesis of distributed reactive control protocols for a Boolean network in a distributed manner. We start with a directed acyclic graph representing a network of Boolean subsystems and a global contract, given as an assumption-guarantee pair. Assumption captures the environment behavior, and guarantee is the requirements to be satisfied by the system. Local assumption-guarantee contracts, together with local control protocols ensuring these local contracts, are computed recursively for each subsystem based on the partial order structure induced by the directed acyclic graph. By construction, implementing these local control protocols together guarantees the satisfaction of the global assumption-guarantee contract. Moreover, local control protocol synthesis reduces to quantified satisfiability (QSAT) problems in this setting. We also discuss structural properties of the network that affect the completeness of the proposed algorithm. As an application, we show how an aircraft electric power system can be represented as a Boolean network, and we synthesize distributed control protocols from a global assumption-guarantee contract. The assumptions capture possible failures of the system components, and the guarantees capture safety requirements related to power distribution.**

I. Introduction

Cyber-physical systems should be able to react to internal and external events since they are expected to work in dynamic settings. Ability to algorithmically synthesize reactive control protocols that can guarantee correct and safe behavior of such systems, can reduce the cost and efforts spent in extensive simulation and testing phase. This fact motivated the research on correct-by-construction control synthesis approaches that start with accurate system models and formal system requirements and that algorithmically construct the controllers to achieve desired closed loop behavior [19]. However, similar principled approaches for synthesis of distributed controllers are scarce partly due to the fact that general distributed synthesis problem is computationally hard [16] when temporal logic requirements are considered.

There is a relation between the control architecture and difficulty of the distributed synthesis problem [9]. Recently some fragments of linear temporal logic (LTL) [6] and some control architectures [12] that render the synthesis problem decidable have been identified. However, efficient algorithms for the decidable fragments are somehow lacking [6]. In this work, we consider a simplified version of the distributed control synthesis problem where the system is modeled as a Boolean network and specifications are given as propositional formulae on environment and system variables. This is a decidable problem since the set of all the controllers is enumerable. It can also be seen as a very special case of distributed LTL synthesis with invariance specifications and memoryless system models. We start by showing that the problem is generally hard (i.e., NEXPTIME-complete) even in this simplified setting. We then propose a sound algorithm to synthesize distributed controllers in a relatively efficient way. The algorithm is shown to be complete when the specification and the control architecture satisfy certain structural properties.

Boolean networks are introduced in [11] to model and analyze the gene circuits. Recent results on control and analysis of such networks can be found in [7]. However, to the best of our knowledge, distributed control problem has not been addressed before in this domain. Our algorithm is motivated by interface theories for stateless components [8] and contract-based verification techniques [2], [4]. However, instead of verification, we consider controllable components and synthesis of controllers. In particular, the proposed algorithm starts from a global assume-guarantee specification and maps it to local assume-guarantee specifications for individual subsystems. Finding these local assume-guarantee specifications leads to a modular design framework. For instance, local controllers for subsystems can be synthesized independently via SAT or a local subsystem can be replaced with another subsystem that satisfies the same local assume-guarantee specifications.

In the second part of the paper, we consider an application of the proposed approach to the co-synthesis of a distributed control architecture and correct-by-construction control protocols for an aircraft electric power system (EPS) [14]. We show how the steady state behavior of an electric power system circuit can be modeled as a Boolean network and synthesize distributed controllers using the proposed algorithm.

II. Preliminaries

A. Notation

The set of positive integers up to and including n is denoted by $\mathbb{N}_n$. For a set K, its cardinality is denoted by $|K|$, and its complement is denoted by K^c. The Boolean domain $\{True, False\}$ is denoted by $\mathbb{B}$. Given a set X of Boolean variables, the set of all valuations is denoted by

A preliminary version of this work is presented in [18] as a work-in-progress abstract. This is an extended version of the paper [17] to appear in the Proceedings of the IEEE Multi-Conference on Systems and Control (MSC), Buenos Aires, Argentina, September 19-22, 2016. This work is supported in part by NSF grants CNS-1446298 and ECCS-1553873, and DARPA grant N66001-14-1-4045.

The authors are with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109, email: ysahin,necmiye@umich.edu.

$V_x = \mathbb{B}^{|X|}$ and called the domain of X, and a specific valuation is denoted by $x \in V_x$.

B. Graph theory

A *graph* is a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a set of nodes (vertices) and $\mathcal{E} \subset (\mathcal{V} \times \mathcal{V})$ is a set of edges.

The graph $\mathcal{G}$ is called *directed* (digraph) when edges have a direction. We say that edge $e = (u, v) \in \mathcal{E}$ points from u to v. A *path* σ on a digraph is a sequence $\sigma = \{e_1, \ldots, e_k\}$ of consecutive edges $e_i = (v_s^i, v_d^i)$, i.e., with $v_d^i = v_s^{i+1}$ for all $1 \leq i < k$. The path σ is said to be a path from node v_s^1 to node v_d^k. A path with identical start and end nodes (i.e., $v_s^1 = v_d^k$) is called a *cycle*. If directed graph $\mathcal{G}$ has no cycles, it is called a *directed acyclic graph*. Furthermore a directed acyclic graph is called a *tree* if any two vertices are connected with at most one path. Any disjoint union of trees is called a *forest*. A node without any outgoing edge is called a *leaf node*.

A graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ is called an *induced subgraph* of $\mathcal{G}$ if $\mathcal{V}' \subseteq \mathcal{V}$ and $\mathcal{E}' \subset ((\mathcal{V}' \times \mathcal{V}') \cap \mathcal{E})$. All the subgraphs we consider in this paper are induced subgraphs, so, we call them just subgraph for short. A digraph is called *strongly connected* if there is a path from each node to every other node. The *strongly connected components* of a directed graph $\mathcal{G}$ are its maximal strongly connected subgraphs. Here maximal is used in the sense that no strongly connected component of $\mathcal{G}$ is a subgraph of some other strongly connected subgraph of $\mathcal{G}$.

C. Boolean systems and networks

A *Boolean system* S is a tuple of the form $\langle U, E, Y, f \rangle$ where

- $U = \{u^{(1)}, \ldots, u^{(n_u)}\}$ is the set of Boolean control inputs with domain $V_u \doteq \mathbb{B}^{n_u}$,
- $E = \{e^{(1)}, \ldots, e^{(n_e)}\}$ is the set of Boolean environment (uncontrolled) inputs, disjoint from U, with domain $V_e \doteq \mathbb{B}^{n_e}$,
- $Y = \{y^{(1)}, \ldots, y^{(n_y)}\}$ is the set of Boolean outputs with domain $V_y \doteq \mathbb{B}^{n_y}$, and
- $f : V_u \times V_e \to V_y$ is the system function, that maps the inputs to the outputs.

The j^{th} component of the system function is denoted by $f^{(j)} : V_u \times V_e \to V_{y^{(j)}}$. Note that this system definition corresponds to a memoryless or stateless system, that is,

$$y = f(u, e) \tag{1}$$

where $y \in V_y$, $u \in V_u$ and $e \in V_e$ are valuations of the outputs, control inputs and environment inputs, respectively.

Given two systems $S_1 = \langle U_1, E_1, Y_1, f_1 \rangle$ and $S_2 = \langle U_2, E_2, Y_2, f_2 \rangle$, a *serial interconnection* from S_1 to S_2 is formed by equating a set of outputs of S_1 to a set of environment inputs of S_2. We denote the set of shared variables in a serial interconnection as

$$I_{1,2} = \{(k, l) \mid y_1^{(k)} = e_2^{(l)}\}, \tag{2}$$

and say S_1 is connected to S_2 when the set of shared variables is nonempty, i.e., $I_{1,2} \neq \emptyset$.

A *Boolean network* is a tuple $S = \langle \{S_i\}_{i=1}^n, \mathcal{I} \rangle$ that consists of a collection of subsystems $S_i = \langle U_i, E_i, Y_i, f_i \rangle$, for $i \in \mathbb{N}_n$, and an interconnection structure $\mathcal{I} = \{I_{i,j}\}_{i,j \in \mathbb{N}_n}$, where $I_{i,j}$ is defined as in (2). The interconnection structure represents which subsystems are connected to which others through which variables. The interconnection structure induces a digraph $\mathcal{G}_S = (\mathcal{V}_S, \mathcal{E}_S)$ called the *system graph*, where $\mathcal{V}_S = \{S_1, \ldots, S_n\}$ and $\mathcal{E}_S = \{(S_i, S_j) \mid I_{i,j} \neq \emptyset\}$. The set $E_i^{int} \subseteq E_i$ of internal inputs for a system S_i is defined as $E_i^{int} = \{e_i^{(l)} \in E_i \mid \exists j \in \mathbb{N}_n, k \in \mathbb{N}_{n_{y_j}}, (k, l) \in I_{j,i}\}$. Remaining inputs of S_i are called external inputs and defined as $E_i^{ext} = E_i \setminus E_i^{int}$.

For well-posedness of the network, we assume the following: (i) If two systems are connected to a third one, they are connected through different environment inputs. That is, if there exist some $i, j \in \mathbb{N}_n$, $i \neq j$ with $(k_1, l) \in I_{i,j}$, then there does not exist any k_2 with $(k_2, l) \in I_{i',j}$ for any $i' \neq i$. (ii) The system graph $\mathcal{G}_S$ is a directed acyclic graph. Note that in the memoryless Boolean setting, cycles in the interconnection structure lead to mostly not well-defined algebraic loops, therefore such interconnections are not considered. In fact, the Boolean network S itself is a Boolean system whose inputs, outputs and system function can be derived from those of its subsystems, and the interconnection structure.

D. Specifications

We consider specifications in the form of *assumption-guarantee* pairs given in terms of Boolean-valued functions. This section introduces some terminology regarding Boolean-valued functions and assume-guarantee specifications.

Let $\varphi : V_x \to \mathbb{B}$ be a Boolean-valued function of variables in X. Each φ can be equivalently represented by its satisfying set $[\![\varphi]\!] \subseteq V_x$ defined as follows:

$$[\![\varphi]\!] \doteq \{x \in V_x \mid \varphi(x) = True\}. \tag{3}$$

Propositional operations on Boolean-valued functions are defined in the usual way. The symbols $\neg, \wedge$ and $\vee$ are used for logical operations *NOT* (negation), *AND* (conjunction) and *OR* (disjunction), respectively and performs as follows: $[\![\neg\psi]\!] \doteq [\![\psi]\!]^c$, $[\![\psi_1 \wedge \psi_2]\!] \doteq [\![\psi_1]\!] \cap [\![\psi_2]\!]$, and $[\![\psi_1 \vee \psi_2]\!] \doteq [\![\psi_1]\!] \cup [\![\psi_2]\!]$.

Additional operations such as *XOR* (exclusive-or) denoted by $\oplus$, can be defined using the main operators above: $\psi_1 \oplus \psi_2 \doteq (\psi_1 \wedge \neg\psi_2) \vee (\neg\psi_1 \wedge \psi_2)$.

Definition 1: Let $\varphi : V \to \mathbb{B}$ be a Boolean-valued function and assume $V = \prod_{i \in J} V_i$. The *projection* of φ onto a set $I \subseteq J$ of variables is another Boolean function $\varphi|_I : \prod_{i \in I} V_i \to \mathbb{B}$ whose satisfying set is defined as

$$[\![\varphi|_I]\!] \doteq \left\{ x \in \prod_{i \in I} V_i \mid \exists y \in \prod_{j \notin I} V_j \text{ such that } (x, y) \in [\![\varphi]\!] \right\}.$$

When the identity of the variables are clear from the context, the order of the tuple is ignored. In other words, saying $(x, y) \in [\![\varphi]\!]$ is equivalent to saying $(y, x) \in [\![\varphi]\!]$ and vice versa.

Assumptions and guarantees are Boolean functions that capture the *a priori* knowledge about the uncontrolled inputs and the desired safe behavior of the system outputs, respectively.

Definition 2: An *assumption* $A : V_e \to \mathbb{B}$ for a Boolean system S is a Boolean-valued function of its uncontrolled inputs.

We say that $(e_1^{ext}, \ldots, e_i^{ext}, \ldots, e_n^{ext}) \in V_e$ is *admissible* if it is in the satisfying set of A. With a slight abuse of terminology, we also say e_i^{ext} is admissible, for the sake of convenience.

Definition 3: A *guarantee* $G : V_y \to \mathbb{B}$ for a Boolean system S is a Boolean-valued function of its outputs.

For a Boolean network S with subsystems $S_1, \ldots, S_n$, by convention, A denotes an assumption with domain $V_e = \prod_{j=1}^{n} V_{e_j^{ext}}$. Furthermore, $A^{\downarrow(i)}$ denotes an assumption associated with subsystem S_i with domain $V_{e_i^{ext}}$. Note that for a given assumption A, its projection $A|_{E_i^{ext}}$ can always be denoted by $A^{\downarrow(i)}$ as it only contains variables from E_i. Similarly, G denotes a guarantee with domain $V_y = \prod_{j=1}^{n} V_{y_j}$ and $G^{\downarrow(i)}$ denotes a guarantee associated with subsystem S_i with domain V_{y_i}.

Definition 4: A formula of the form

$$\varphi \doteq \bigwedge_{k=1}^{n_c} (A_k \to G_k) \tag{4}$$

is called a *global contract* and denoted by $\mathbb{C} \doteq \{[A_k, G_k]\}_{k=1}^{n_c}$.

Definition 5: A formula of the form

$$\varphi^{(i)} \doteq \bigwedge_{k=1}^{n_c} (A_k^{\downarrow(i)} \to G_k^{\downarrow(i)}) \tag{5}$$

is called a *local contract* for S_i and denoted by $\mathbb{C}_i \doteq \{[A_k^{(i)}, G_k^{\downarrow(i)}]\}_{k=1}^{n_c}$.

E. Control protocol synthesis

Given a system S, a *control protocol* $\pi : V_e \to V_u$ maps the environment inputs to control inputs. When a control protocol π is implemented on a system S, *controlled system* is governed by the following input-output relation:

$$y = f(\pi(e), e), \tag{6}$$

since $u = \pi(e)$.

Given an assumption-guarantee pair, the control synthesis problem aims to find a protocol π that sets the control inputs such that $A \to G$ is satisfied. For any $e \in [\![A]\!]$, determining if there exist a $u \in V_u$ so that G evaluates $True$, is a quantified Boolean SAT problem:

$$\forall e \in V_e : \exists u \in V_u : A(e) \to G(y) \tag{7}$$

where y is given as in the Equation (1). If the quantified SAT problem above can be solved, then the set of all solutions can be used as the control protocol. In this case, we say $A \to G$ is *realizable*.

III. Problem Statement

We are interested in synthesizing distributed controllers for a Boolean network to achieve a common goal while each local controller has access only to its own inputs. The general form of the problem can be formally stated as follows:

Problem 1: Given a Boolean network $S = \langle \{S_i\}_{i=1}^{n}, \mathcal{I} \rangle$ and a Boolean formula φ over the input and output variables $\bigcup_{i=1}^{n}(E_i \cup Y_i)$, find local controllers $\pi_i : V_{e_i} \to V_{u_i}$ such that when these local controllers are implemented together, the controlled system satisfies φ for all possible inputs.

Next, we show that even in this relatively simple Boolean setting, the distributed synthesis problem is quite hard. The main difficulty arises from the fact that each local controller only has partial information to base their decisions on. In order to show the hardness of the problem, we reduce dependency quantifier boolean formula game (DQBFG) to Problem 1. Let us first introduce DQBFG [10].

DQBFG is a three player game, consisting of players B (black), $W1$ (white 1) and $W2$ (white 2). An instance of DQBFG is a Boolean formula φ in variables $X_1 \cup X_2 \cup Z_1 \cup Z_2$. For each $i \in \{1, 2\}$, player Wi has access only to the variables in $X_i \cup Z_i$. First, player B chooses an assignment for variables in $X_1 \cup X_2$. Then, player $W1$ chooses an assignment for variables in Z_1 and player $W2$ chooses an assignment for variables in Z_2. Note that the order of the decisions made by $W1$ and $W2$ does not matter since they do not see each other's decisions. White team wins if φ is true in the end. The decision problem DQBFG is whether a winning strategy for the white team in this game exists. DQBFG belongs to the complexity class NEXPTIME-complete [10].

Theorem 1: The decision version of Problem 1 is NEXPTIME-complete.

Proof: We prove our claim by showing that DQBFG is a special case of our problem. We can imagine environment taking over the role of B in our setting. Boolean network S consists of two subsystems S_1, S_2 associated with $W1$ and $W2$ respectively, where

- $\mathcal{I} \doteq \{I_{1,1}, I_{1,2}, I_{2,1}, I_{2,2}\}$ with $I_{i,j} \doteq \emptyset$ for all $i, j = 1, 2$,
- $E_i \doteq X_i$ and $U_i \doteq Z_i$ for $i = 1, 2$,
- $Y_i \doteq E_i \cup U_i = X_i \cup Z_i$ and the system function is the identity map for each subsystem, i.e., $f_i : (u_i, e_i) \mapsto (u_i, e_i)$ for $i = 1, 2$.

The Boolean formula φ' is obtained by replacing each variable X_i with Z_i in φ with the corresponding variables in Y_i. Then, the decision version of Problem 1, that is, verifying the existence of local controllers to satisfy a Boolean function φ' for the network S as defined above, is equivalent to DQBFG. This proves NEXPTIME-hardness.

Now we show that Problem 1 is in NEXPTIME. A naive algorithm for Problem 1 that enumerates all possible local controller tuples π^i for $i = 1, \ldots, n$ and checks the satisfaction of φ runs in NEXPTIME since there are $\mathcal{O}(2^{\sum_{i \in \mathbb{N}_n} |U_i| 2^{|E_i|}})$ local controller tuples. Thus the complexity of the decision version of Problem 1 is NEXPTIME-complete. ∎

Given the hardness result above, we seek to break the problem into more manageable pieces. We assume that φ is given in the form of a global contract as in (4). Note that this is without loss of generality as any φ can be written in this form. We consider the following variant of the problem.

Problem 2: Given a Boolean network $S = \langle \{S_i\}_{i=1}^n, \mathcal{I} \rangle$, and a global contract $\mathbb{C} = \{[A_k, G_k]\}_{k=1}^{n_c}$, for $i = 1, \ldots, n$, find local contracts of the form $\mathbb{C}_i = \{[A_k^{\downarrow(i)}, G_k^{\downarrow(i)}]\}_{k=1}^{n_c(i)}$ and local controllers $\pi_i : V_{e_i} \to V_{u_i}$, satisfying these local contracts, such that when these local controllers are implemented together, the controlled system satisfies $\mathbb{C}$ for all possible inputs.

Note that Problem 2 is essentially equivalent to Problem 1. However, if the structure of the problem allows the local contracts $\mathbb{C}_i$ to be computed efficiently, then obtaining a local controller from a local contract is a full information synthesis problem that can be solved via QSAT. Complexity of QSAT is PSPACE-complete, which is lower in the complexity hierarchy than NEXPTIME-complete [15] . Moreover, there are off-the-shelf highly optimized tools for solving QSAT problems making their solution practically feasible. In what follows, we provide a sound algorithm to solve Problem 2. Some structural properties of $\mathbb{C}$ and the interconnection structure of S that render the algorithm complete are also discussed.

IV. Algorithm

In this section, we impose a certain structure on the specification φ so that Problem 2 can be reduced to a number of QSAT problems. In particular, we let the specification φ be given by a single assumption-guarantee pair, $\mathbb{C} = [A, G]$ and propose Algorithm 1 to solve for this case.

An overview of the algorithm is as follows. Let S_i be a subsystem that is a leaf node of the system graph. We first compute a local assumption $A^{\downarrow(i)}$ (line 3) and a guarantee $G^{\downarrow(i)}$ (line 4). Then, we synthesize the local controller, which is a full information synthesis problem that can be solved by QSAT. If the local assumption-guarantee pair is unsatisfiable, we constrain the internal inputs of S_i to make it realizable using the so called *least restrictive assumptions* (line 6). Since internal inputs of S_i are outputs of other subsystems, this assumption becomes a guarantee the rest of the systems must fulfill (line 8). Once we synthesize the controller, we discard S_i from the system graph and we pick another leaf node S_j and continue in the same manner until every subsystem has a local controller. Different subroutines of the algorithm are explained in detail next.

A. Finding local assumptions

Local assumptions are simply found by projection as shown in line 3 the Algorithm 1. By definition of the projection operator, we have $A^{\downarrow(i)}$ depend only on local external variables. Moreover, projection ensures that local assumptions do not restrict the environment more than the global assumption and does so in the "best" possible way. In particular, we have the following property that will be useful in proving the soundness of the algorithm.

Algorithm 1 Distributed_Synthesis($S, \mathbb{C}$)

Input: A Boolean network S and global contract $\mathbb{C} = [A, G]$

1: **if** S is empty **then return** $True$
2: Let S_j be a leaf node of $\mathcal{G}_S$
3: $A^{\downarrow(j)} = A|_{E_j^{ext}}$
4: Γ = distribute_guarantee(G, S_j)
5: **for** $\gamma \in \Gamma$ **do**
6: $(\lambda_{lra}^{(j)}, \pi_j)$ = find_lra($S_j, A^{\downarrow(j)}, G^{\downarrow(j)}$)
7: **if** $\lambda_{lra}^{\downarrow(j)}$ is $False$ **then return** $False$
8: $\mathbb{C}' = \left[A, \left(G^{\uparrow(j)} \wedge \lambda_{lra}^{(j)}\right)\right]$
9: S' = delete(S, S_j)
10: flag = Distributed_Synthesis($S', \mathbb{C}'$)
11: **if** flag **then return** $True$
12: **return** $False$

Proposition 1: Let $A : \prod_{i=1}^n V_{e_i^{ext}} \to \mathbb{B}$ be the global assumption. Define S_i's local assumption $A^{\downarrow(i)} \doteq A|_{V_{e_i^{ext}}}$ for $i \in \mathbb{N}_n$. Then local assumptions are less restrictive than the global assumption, i.e.,

$$[\![A]\!] \subset \prod_{i=1}^{n} [\![A^{\downarrow(i)}]\!]. \tag{8}$$

Moreover, (8) fails to hold if we replace any $A^{\downarrow(i)}$ with $\bar{A}^{\downarrow(i)}$ where $[\![\bar{A}^{\downarrow(i)}]\!] \subset [\![A^{\downarrow(i)}]\!]$.

Proof: Assume that $(e_1^{ext}, \ldots, e_n^{ext}) \in [\![A]\!]$. By construction, $e_i^{ext} \in [\![A^{\downarrow(i)}]\!]$ for all $i \in \mathbb{N}_n$. Then $(e_1^{ext}, \ldots, \ldots, e_n^{ext}) \in \prod_{i=1}^n [\![A^{\downarrow(i)}]\!]$. This implies (8).

For any arbitrary $i \in \mathbb{N}_n$, let $\bar{A}^{\downarrow(i)}$ be a Boolean-valued function satisfying $[\![\bar{A}^{\downarrow(i)}]\!] \subset [\![A^{\downarrow(i)}]\!]$ and $[\![\bar{A}^{\downarrow(j)}]\!] = [\![A^{\downarrow(j)}]\!]$ for every other $j \neq i$. Now let $e_i^{ext} \in [\![A^{\downarrow(i)}]\!] \setminus [\![\bar{A}^{\downarrow(i)}]\!]$. By construction of $A^{\downarrow(i)}$, there exists an environment valuation $(e_1^{ext}, \ldots, e_i^{ext}, \ldots, e_n^{ext}) \in [\![A]\!]$. However, $(e_1^{ext}, \ldots, e_i^{ext}, \ldots, e_n^{ext}) \in \prod_{j=1}^n [\![\bar{A}^{\downarrow(j)}]\!]$. Thus (8) is no longer true. ■

B. Finding local guarantees

As opposed to the local assumptions, local guarantees $G^{\downarrow(i)}$ cannot be less restrictive than the global guarantee. No communication is assumed between subsystems, hence any possible combination of outputs allowed by local guarantees should be in the satisfying set of global guarantee, i.e.,

$$\forall i \in \mathbb{N}_n : \forall y_i \in [\![G^{\downarrow(i)}]\!] : (y_1, \ldots, y_n) \in [\![G]\!]. \tag{9}$$

Due to the dependence, local guarantees cannot be computed independently. Our algorithm proceeds by selecting a local guarantee for a subsystem at each iteration, using the notion of a distribution, which is introduced next.

Definition 6: Let $G : \prod_{j=1}^n V_{y_i} \to \mathbb{B}$ be the global guarantee. A *distribution* of G between S_i and the rest of the subsystems is a pair of Boolean valued functions $G^{\downarrow(i)} : V_{y_i} \to \mathbb{B}$ and $G^{\uparrow(i)} : \prod_{j \neq i} V_{y_j} \to \mathbb{B}$ where

$$\forall y_i \in [\![G^{\downarrow(i)}]\!] : \forall y \in [\![G^{\uparrow(i)}]\!] : (y_i, y) \in [\![G]\!]. \tag{10}$$

We denote a distribution as $\gamma = \{G^{\downarrow(i)}, G^{\uparrow(i)}\}$.

Unfortunately, in general, there is no unique "best" way of generating local guarantees, therefore distributions are

not unique. Note that $[\![G^{\downarrow(i)}]\!] \times [\![G^{\uparrow(i)}]\!] \subset [\![G]\!]$ is trivial by Eq. (10). This implies that local guarantees are conservative and they under-approximate $[\![G]\!]$. However, we do not want to restrict the system more than necessary. That is why we compute only the *maximal distributions*. A distribution is called *maximal* when there is no other distribution $\bar{\gamma} = \{\bar{G}^{\downarrow(i)}, \bar{G}^{\uparrow(i)}\}$ that satisfy $[\![G^{\downarrow(i)}]\!] \subset [\![\bar{G}^{\downarrow(i)}]\!]$ and/or $[\![G^{\uparrow(i)}]\!] \subset [\![\bar{G}^{\uparrow(i)}]\!]$. We denote the set of all maximal distributions with $\Gamma = \{\gamma_k\}_k$, which is computed in line 4 of the algorithm. We refer the reader to the appendix for more details on how to compute these distributions.

C. Least restrictive assumptions and controller synthesis

Controller synthesis for any subsystem is essentially a quantified SAT problem as stated in (7). However, if we let internal inputs to take any possible value, then it might not be possible to find an input u_i that renders $G^{\downarrow(i)} = True$. When this is the case, we restrict the internal inputs, which are controlled by S_i's ancestors, to a certain set to achieve $G^{\downarrow(i)}$. While doing so, we would like to be as permissive as possible.

Definition 7: Given a local contract $\mathbb{C}_i = [A^{\downarrow(i)}, G^{\downarrow(i)}]$, the set of all internal inputs that makes the contract realizable is called the *least restrictive assumption*. It is denoted with $\lambda^{(i)}_{lra}$ where

$$[\![\lambda^{(i)}_{lra}]\!] \doteq \{e^{int}_i \in E^{int}_i \mid A^{\downarrow(i)} \rightarrow G^{\downarrow(i)} \text{ is realizable}\}. \quad (11)$$

In other words, the least restrictive assumption gives the set of internal inputs that makes the guarantee realizable. Any internal input outside of this set makes the guarantee unsatisfiable.

After computing the least restrictive assumption, we update the local contract as

$$\mathbb{C}_i = \left[\left(A^{\downarrow(i)} \wedge \lambda^{\downarrow(i)}_{lra}\right), G^{\downarrow(i)}\right]. \quad (12)$$

Note that by definition of least restrictive assumption, $\mathbb{C}_i$ is realizable. Having a realizable local contract, the control protocol is synthesized by solving the respective QSAT problem. Line 6 in Algorithm 1 performs these two operations simultaneously.

On the other hand, the least restrictive assumption $\lambda^{(i)}_{lra}$ imposes new guarantees to the ancestors of S_i. We change internal inputs with their output correspondents by examining the interconnection structure $\mathcal{I}$. Finally we update the global contract for the remaining subsystems as

$$\mathbb{C}' = [A, (G^{\uparrow(i)} \wedge \lambda^{(i)}_{lra})]. \quad (13)$$

D. Algorithm Analysis

In this section we show that solutions returned by Alg. 1 are correct. Then we introduce conditions on the system graph and the specifications that renders the algorithm complete. Finally we discuss the complexity of the proposed method.

Theorem 2 (Soundness): For arbitrary Boolean network S, if the specification is given with $\mathbb{C} = [A, G]$, then Alg. 1 is sound.

Proof: It is enough to show that local contracts results in weaker assumptions on environment inputs and stronger restrictions on system outputs. It can be shown using Boolean algebra that

$$\left(\bigwedge_{i=1}^{n}\left(A^{\downarrow(i)} \rightarrow G^{\downarrow(i)}\right)\right) \rightarrow \left(\bigwedge_{i=1}^{n} A^{\downarrow(i)} \rightarrow \bigwedge_{i=1}^{n} G^{\downarrow(i)}\right) \quad (14)$$

is a tautology. Also from Eq. (8) and Eq.(9), we can show that $A \rightarrow \bigwedge_{i=1}^{n} A^{\downarrow(i)}$ and $\bigwedge_{i=1}^{n} G^{\downarrow(i)} \rightarrow G$ are tautologies. Therefore,

$$\left(\bigwedge_{i=1}^{n} A^{\downarrow(i)} \rightarrow \bigwedge_{i=1}^{n} G^{\downarrow(i)}\right) \rightarrow (A \rightarrow G) \quad (15)$$

is trivially true.

Also note that the least restrictive assumptions computed for each subsystem are not restrictions on environment. In fact they restrict the outputs, which only strengthens result presented above. Hence, if the algorithm returns a controller that satisfies the local contracts, the global contract is satisfied. ■

The completeness requires additional assumptions on the system graph and the specification.

Theorem 3 (Completeness): Let the specifications given with the contract $\mathbb{C} = [A, G]$ and the system graph induced by the Boolean system S be $\mathcal{G}_S$. Alg. 1 is complete if

1) $A = \bigwedge_i A^{\downarrow(i)}$,
2) $G = \bigwedge_i G^{\downarrow(i)}$, and
3) $\mathcal{G}_S$ is a forest.

In other words the specification is given as

$$\mathbb{C} = \left[\bigwedge_i A^{\downarrow(i)}, \bigwedge_i G^{\downarrow(i)}\right]. \quad (16)$$

Proof: We prove by induction that if Alg. 1 fails to return a protocol then there does not exist one.

When the number of subsystems is one, centralized and distributed algorithms are identical. If the respective QSAT problem is not satisfiable, then there does not exist any control protocol to achieve the task.

Assume that the Alg. 1 is complete for an arbitrary forest $\bar{\mathcal{G}}_S$ with n nodes and any specification $\bar{\varphi}$ given in the form of Eq. (16).

Also let $\mathcal{G}_S$ be an arbitrary forest with $(n+1)$ nodes and let

$$\mathbb{C} = \left[\bigwedge_{i=1}^{n+1} A^{\downarrow(i)}, \bigwedge_{i=1}^{n+1} G^{\downarrow(i)}\right] \quad (17)$$

be its specification. Without loss of generality, assume that $S_{(n+1)}$ is a leaf node and the local contract that is computed according to Alg. 1 is given as

$$\mathbb{C}_{(n+1)} = \left[\left(A|_{V_{E^{ext}_{n+1}}} \wedge \lambda^{(n+1)}_{lra}\right), G^{\downarrow(n+1)}\right].$$

First assume that $\lambda^{(n+1)}$ is $False$ and the local contract $\mathbb{C}_{(n+1)}$ is not satisfiable. This means that there exists at least one admissible environment valuation $e^{ext}_{(n+1)}$ such that no matter what the internal inputs are, $G^{\downarrow(n+1)}$ is unsatisfiable.

Since the environment inputs are uncontrolled, no control protocol can overcome this problem. Thus no distributed (or central) control protocol exists for the given system and specifications.

Now assume that $\lambda^{(n+1)}$ is not $False$ and the local contract $\mathbb{C}_{(n+1)}$ is satisfiable. Then the global contract is updated as

$$\mathbb{C}' = \left[A, \left(G_{lra}^{\uparrow(n+1)} \wedge \lambda_{lra}^{\downarrow(n+1)}\right)\right]. \tag{18}$$

Note that distribution is unique and $G_{lra}^{\uparrow(n+1)} = \bigwedge_{i\neq n+1} G^{\downarrow(i)}$. Let S_j denote the parent of S_i. Then $\lambda_{lra}^{(n+1)}$ is an additional guarantee that involves variables only from S_j. Now define $(G')^{\downarrow(i)} = G^{\downarrow(i)}$ for $i \neq j$ and $(G')^{\downarrow(j)} = G^{\downarrow(j)} \wedge \lambda_{lra}^{(n+1)}$. Then we can write

$$\mathbb{C}' = \left[\bigwedge_{i=1}^{n} (A')^{\downarrow(i)}, \bigwedge_{i=1}^{n} (G')^{\downarrow(i)}\right] \tag{19}$$

where $A' = A|_{\prod_{j\neq i}^{n}}$. Now we are left with an arbitrary forest with n nodes and a specification given in the form of Eq. (16). Thus the Alg. 1 is complete. ■

Remark 1: A few remarks on the complexity of the algorithm are in order. In the worst case, the proposed algorithm requires solving $\mathcal{O}(n2^{\Sigma_i|Y_i|+\max_i|E_i|})$ SAT problems each with complexity $\mathcal{O}(2^{\max_i|U_i|})$, where n is the number of subsystems. Under the assumptions of Theorem 3, it reduces to $\mathcal{O}(n2^{\max_i|E_i|})$ since the distribution can simply be computed using projection and is unique. Recall that the complexity of the general problem is $\mathcal{O}(2^{\sum_{i\in\mathbb{N}_n}|U_i|2^{|E_i|}})$; therefore with the assumed structure, one of the exponents is eliminated.

V. Illustrative Examples

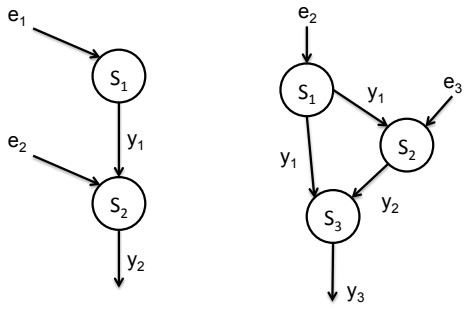

Fig. 1. (Left): A system graph satisfying the forest condition. (Right): System graph where forest condition is violated.

In this section we present examples to illustrate the proposed method. We use Algorithm 1 to solve Example 1 which satisfies all conditions for completeness.

Example 1: For the system graph seen in Fig. 1 (Left), let the global contract be $\mathbb{C} = [e_1, y_2]$, where

$$f_1 : (e_1, u_1) \mapsto u_1$$
$$f_2 : (e_2, y_1, u_2) \mapsto ((e_2 \vee y_1) \wedge u_2).$$

In order to synthesize local controllers, we start by finding the local assumption for S_2. Using Boolean projection, it can be found as

$$A^{\downarrow(2)} = A|_{V_{e_2^{ext}}} = True.$$

Note that e_2 can take any possible value because the global contract does not put any restrictions on S_2's environment inputs. This implies that the local assumption is trivially true. Then we distribute the guarantee $G = y_2$ and compute the distribution. In this example guarantee depends only on S_2's inputs, thus the solution is unique and it can be written as

$$G^{\downarrow(2)} = y_2, \quad G^{\uparrow(2)} = True.$$

Having local assumption-guarantee pair, we compute the least restrictive assumption as $\lambda_{lra}^{(2)} = y_1$.

By definition, as long as $\lambda_{lra}^{(2)}$ is satisfied by its ancestors, S_2 can satisfy its guarantee under all admissible environment conditions. Note that $\lambda_{lra}^{(2)}$ introduces an additional guarantee for the remaining subsystems, i.e. S_1.

Upon synthesizing a local controller for S_2, we update the global contract $\mathbb{C}' = [e_1, y_1]$.

Now we delete S_2 from the system graph and continue with S_1. First, the local assumption and guarantee are computed:

$$A^{\downarrow(1)} = A|_{V_{e_1^{ext}}} = e_1 \qquad G^{\downarrow(2)} = \lambda_{lra}^{(2)} = y_1$$

Note that the local contract is satisfiable. Thus we can synthesize a local controller for S_1 by solving the respective QSAT problem.

Now we provide three examples and show why we need the conditions on specifications and system graph. To begin with, when the first condition in Theorem 3 is satisfied, (8) becomes an equality.

In other words, Boolean projection does not result in information loss. When this condition is violated, some information is lost. This might lead to incomplete solutions as illustrated by the next example.

Example 2: For the system graph seen Fig. 1 (Left), let $\mathbb{C} = [e_1 \oplus e_2, y_2]$ where

$$f_1 : (u_1, e_1) \mapsto (e_1 \wedge u_1)$$
$$f_2 : (u_1, e_2, y_1) \mapsto ((e_2 \vee y_1) \wedge u_2).$$

In this example, the assumption cannot be written as a conjunction of two local assumptions, hence the completeness conditions are violated. Let us apply the algorithm to see why this causes a problem. We start by finding local assumption-guarantee pair and the corresponding least restrictive assumption.

$$A^{\downarrow(2)} = A|_{V_{e_2^{ext}}} = True, \qquad G^{\downarrow(2)} = y_2,$$
$$G^{\uparrow(2)} = True, \qquad \lambda_{lra}^{(2)} = y_1.$$

Then, we update the contract as $\mathbb{C}' = [e1 \oplus e_2, y_1]$. We delete S_2 and move onto S_1. Unfortunately the local assumption $A^{\downarrow(1)} = A|_{V_{e_1^{ext}}} = True$ together with $G^{\downarrow(1)} = y_1$ is unsatisfiable.

On the other hand, the following local contracts leads to local controllers satisfying the global specifications: $\mathbb{C}_2 = [e_2 \vee y_1, y_2]$, $\mathbb{C}_1 = [e_1, y_1]$.

When the second condition in Theorem 3 is violated, $[\![G^{\downarrow(i)}]\!] \times [\![G^{\uparrow(i)}]\!] \subset [\![G]\!]$ is a strict subset relation. This implies that Boolean distribution is conservative. Next example shows how this results in incomplete solutions.

Example 3: For the system graph seen Fig, let $\mathbb{C} = [True, y_1 \vee y_2]$ where

$$f_1 = (u_1, e_1) \mapsto e_1 \wedge u_1$$

$$f_2 = (u_2, e_2, y_1) \mapsto (e_2 \vee u_2) \wedge \neg y_1$$

In this example, second condition is violated. Then the distribution is not unique. We have

$$\gamma_1 = \{G_1^{\downarrow(2)} = y_2, G_1^{\uparrow(2)} = True\}$$

or

$$\gamma_2 = \{G_2^{\downarrow(2)} = True, G_2^{\uparrow(2)} = y_1\}.$$

Using $G_1^{\uparrow(2)}$ results in $\lambda^{(2)} = y_1$. Then the local contract $\mathbb{C}_1 = True, y_1$ is not satisfiable. Similarly the second option also forces S_1 to set $y_1 = True$. Thus again, resulting in an unsatisfiable local contract.

However, the following local contracts can be used to synthesize local controllers satisfying $\mathbb{C}$: $\mathbb{C}_2 = [\neg y_1, y_2]$ and $\mathbb{C}_1 = [e_1, y_1]$.

Finally, if a subsystem has multiple parents, the least restrictive assumption might lead to violation of the second condition. This is demonstrated by the next example.

Example 4: For the system graph seen Fig. 1 (Right), let $\mathbb{C} = [True, y_3]$ where

$$f_1 : (u_1, e_1) \mapsto (e_1 \wedge u_1)$$
$$f_2 : (u_2, e_2, y_1) \mapsto (e_2 \vee u_2) \wedge \neg y_1$$
$$f_3 : (u_3, y_2, y_1) \mapsto (y_1 \vee y_2).$$

In this example, after synthesizing a controller for S_3, global contract is updated to $\mathbb{C}' = [True, y_1 \vee y_2]$.

Note that the remaining problem is equal to Example 3, thus the solution is not complete.

Finally, it is worth mentioning that it is easy to construct examples where there exists a distributed controller but some of the conditions in Theorem 3 fail to hold but the proposed algorithm still finds a valid solution.

VI. Case Study

In this section, we model electric power system of an aircraft as a Boolean network and and translate the specifications into an assumption-guarantee pair. Then we use Algorithm 1 to synthesize distributed controllers. Figure 2 (Left) shows a circuit for power generation and distribution in an aircraft in the form of a single-line diagram [13], a simplified notation for drawing three-phase power systems. To model the system as a Boolean network, we first represent the electric power system topology by a digraph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ that indicates the power flow directions. The set $\mathcal{V}$ of nodes in the graph consists of components such as generators, rectifier units, buses, transformers and dummy nodes where a collection of wires meet. The set $\mathcal{E}$ of edges contains contactors and solid wire links between other components. We locate the maximal connected components of this graph; and treat each connected component as a Boolean subsystem. This results in the interconnection structure shown in Fig. 2 (Right).

In order to analyze the steady-state behavior of the electric power system in different failure modes, we use a discrete abstraction. The status of contactors can either be *open* or *closed*, i.e. Boolean $\{False, True\}$. Contactors are controllable (control inputs). Elements in the sets of generators and rectifier units and transformers are uncontrollable (environment inputs), and can take values of $True$ (i.e., the component is online and outputting the correct voltage), or $False$ (i.e., the component failed, no power output, open circuit). Power availability on a bus (outputs) depends on the status of generators, rectifier units, transformers and contactors. We say that there is a *live path* between two components if there exists a simple path in the graph $\mathcal{G}$ that connects the two nodes corresponding to these components, there is no offline component along the path including end nodes, and the contactors along this path are all closed. The status of a bus B can be (i) $True$ (*powered*): if there is a live path between B and some generator E (not offline by definition of live path), (ii) $False$ (*unpowered*): there is no live path between B and any generator E.

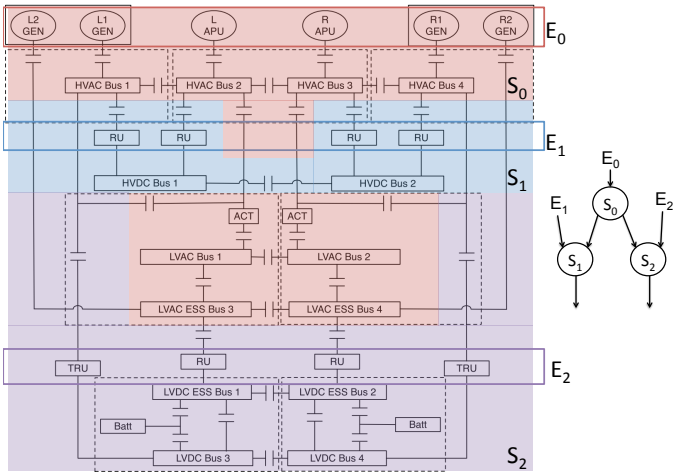

Fig. 2. (Left): Single-line diagram of the aircraft EPS (adapted from [14]). Different colors indicate different subsystems identified. (Right): Corresponding Boolean network.

Boolean network representing the circuit has 22 environment inputs, 32 control inputs and 15 outputs in total. The global assumption, that at least one power source is healthy along with one rectifier from each E_1 and E_2, has 22 variables and 14 conjunctive clauses. We require all buses to be powered and no AC coupling between power sources, that is there should not be a live path between any AC power source. Resulting the global guarantee has 47 variables and 2142 conjunctive clauses. Solving this problem in a centralized manner requires solving a QSAT problem with

69 variables and 2151 clauses. The computation takes 336 seconds using PicoSAT [3] on a laptop with 2.5 GHz Intel Core i7 processor and 16 GB of RAM. Keep in mind that the returned control protocol is a centralized one and is not implementable in a distributed manner in general.

Then we use Alg. 1 to solve for the same case. Now we have three Boolean subsystems. This means we have to solve three QSAT problems that are much smaller with 39, 19, and 24 variables and 485, 369, and 1294 clauses (for S_0, S_1 and S_2, respectively). The system graph is a tree (hence satisfies the forest condition) and the specification satisfies (16). This implies that Alg. 1 is complete. The computation takes 42 seconds on the same laptop and returns a distributed controller.

VII. Conclusion

In this paper, we proposed an algorithm for synthesizing local specifications and corresponding local controllers for a Boolean network to guarantee the satisfaction of a global specification. The algorithm is shown to be sound, and it is complete when the global specification and networks interconnection structure satisfy certain structural properties. An application of the proposed approach is demonstrated by synthesizing distributed controllers for an aircraft electric power system.

In the future, we will extend the presented algorithms to dynamic discrete transition systems and specifications given in certain fragments of linear temporal logic. It would also be interesting to understand the connections between the structural properties identified in this work with those that are known in decentralized control for different system classes and control objectives.

Appendix

Computation of local guarantees

Boolean distribution operation is done by representing the Boolean formulas as graphs and using graph properties. This has two advantages: (i) graphs provide a canonical representation for the satisfying set of the formulas (whereas there could be multiple formulas with the same satisfying set), (ii) once the problem is converted to a graph problem, we leverage well-established algorithms from graph theory to find distributions.

Let us start with some additional graph terminology. Given a graph, a set of nodes, no two of which are connected with an edge is called an *independent set*. A graph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ is called *bipartite* if its nodes can be partitioned into two independent sets, i.e. $\mathcal{V} = \mathcal{V}_1 \cup \mathcal{V}_2$ and $\mathcal{E} \subset \mathcal{V}_1 \times \mathcal{V}_2$. $\mathcal{V}_1$ and $\mathcal{V}_2$ are also called *parts*. If every node in one part is connected to every other node in the other part with an edge, the bipartite graph is called *complete*, i.e., $\mathcal{E} = \mathcal{V}_1 \times \mathcal{V}_2$.

Let $H_c = (\mathcal{V}', \mathcal{E}')$ be a complete bipartite subgraph of a bipartite graph $H = (\mathcal{V}, \mathcal{E})$. Then H_c is called *maximal* if addition of any node to $\mathcal{V}'$ breaks down the completeness. Putting it differently, H_c is not a subgraph of any other H_c' which is also a complete bipartite subgraph of H.

Let S be a Boolean system and $\mathbb{C} = [A, G]$ be its specification. Now assume that S_i is a leaf node in the system graph and we want to compute its local guarantee. In order to perform Boolean distribution, we construct a bipartite graph $H = \left(\left(\prod_{j\neq i} V_{y_j} \bigcup V_{y_i}\right), \mathcal{E}\right)$. In this graph, every node $y_i \in V_{y_i}$ represents a different valuation of S_i's outputs. Similarly each node $y \in \prod_{j\neq i} V_{y_j}$ represent all other outputs valuations. Edges are created between nodes that forms an admissible output, i.e., $\mathcal{E} = \{(y_i, y) \mid (y_i, y) \in [\![G]\!]\}$.

Proposition 2: Let $H_c = \left(\left(\prod_{j\neq i} \bar{V}_{y_j} \bigcup \bar{V}_{y_i}\right), \bar{\mathcal{E}}\right)$ be a subgraph of H where $\bar{V}_{y_k} \subset V_{y_k}$ for all k. Now define $G^{\downarrow(i)} : V_{y_i} \to \mathbb{B}$ and $G^{\uparrow(i)} : \prod_{j\neq i} V_{y_j} \to \mathbb{B}$ with their satisfying set $[\![G^{\downarrow(i)}]\!] = \bar{V}_{y_i}$ and $[\![G^{\uparrow(i)}]\!] = \prod_{j\neq i} \bar{V}_{y_j}$. Then $\gamma = \{G^{\downarrow(i)}, G^{\uparrow(i)}\}$ is a maximal distribution if and only if H_c is a maximal complete bipartite graph.

Proof: Let γ be a distribution. Now define $H_c = \left(\left(\prod_{j\neq i} \bar{V}_{y_j} \bigcup \bar{V}_{y_i}\right), \bar{\mathcal{E}}\right)$ where $\bar{V}_{y_i} = [\![G^{\downarrow(i)}]\!], \prod_{j\neq i} \bar{V}_{y_j} = [\![G^{\uparrow(i)}]\!]$ and $\bar{\mathcal{E}} = \mathcal{E} \cap \left(\prod_{j\neq i} \bar{V}_{y_j} \times \bar{V}_{y_i}\right)$. By definition of distribution $\forall y_i \in [\![G^{\downarrow(i)}]\!] : \forall y \in [\![G^{\uparrow(i)}]\!] : (y_i, y) \in [\![G]\!]$. This implies $\forall y_i \in \bar{V}_{y_i} : \forall y \in \prod_{j\neq i} \bar{V}_{y_j} : (y_i, y) \in \bar{\mathcal{E}}$. Thus H_c is a complete bipartite graph.

Conversely let $H_c = \left(\left(\prod_{j\neq i} \bar{V}_{y_j} \bigcup \bar{V}_{y_i}\right), \bar{\mathcal{E}}\right)$ be a complete bipartite subgraph of H. By definition of complete bipartite graph $\forall y_i \in \bar{V}_{y_i} : \forall y \in \prod_{j\neq i} \bar{V}_{y_j} : (y_i, y) \in \bar{\mathcal{E}} \subset \mathcal{E}$. Now define $G^{\downarrow(i)} : V_{y_i} \to \mathbb{B}$ and $G^{\uparrow(i)} : \prod_{j\neq i} V_{y_j} \to \mathbb{B}$ with their satisfying set $[\![G^{\downarrow(i)}]\!] = \bar{V}_{y_i}$ and $[\![G^{\uparrow(i)}]\!] = \prod_{j\neq i} \bar{V}_{y_j}$. Then $\forall y_i \in [\![G^{\downarrow(i)}]\!] : \forall y \in [\![G^{\uparrow(i)}]\!] : (y_i, y) \in [\![G]\!]$ by construction of H. This implies $\gamma = \{G^{\downarrow(i)}, G^{\uparrow(i)}\}$ is a local guarantee set.

Finally let the distribution $\gamma = \{G^{\downarrow(i)}, G^{\uparrow(i)}\}$ be maximal. This implies that there does not exist another complete bipartite subgraph $\bar{H}_c = \{\prod_{j\neq i} \bar{V}_{y_j} \cup \bar{V}_{y_i}, \bar{\mathcal{E}}\}$ such that $\prod_{j\neq i} \bar{V}_{y_j} \supset \prod_{j\neq i} V_{y_j}$ and $\bar{V}_{y_i} \supset V_{y_i}$. This implies that there does not exist any other subgraph H_c' of H such that H_c is a subgraph of H_c'. Thus H_c is maximal. Going in the other direction is also similar. If H_c is maximal, so is $\gamma = \{G^{\downarrow(i)}, G^{\uparrow(i)}\}$. ∎

Note that with Proposition 2, the problem of finding a distribution reduces to the problem of finding maximal complete bipartite subgraphs. The latter can be done using standard results from graph theory literature [1]. In particular, in our implementation, we have used a variant of the Bron-Kerbosch algorithm [5] to find all maximal distributions.

References

[1] G. Alexe, S. Alexe, Y. Crama, S. Foldes, P. L. Hammer, and B. Simeone. Consensus algorithms for the generation of all maximal bicliques. *Discrete Applied Mathematics*, 145(1):11–21, 2004.

[2] L. Benvenuti, A. Ferrari, E. Mazzi, and A. L. Vincentelli. Contract-based design for computation and verification of a closed-loop hybrid system. In *Proceedings of the Hybrid Systems: Computation and Control*, pages 58–71, 2008.

[3] A. Biere. PicoSAT essentials. *Journal on Satisfiability, Boolean Modeling and Computation (JSAT)*, 2008.

[4] M. Bozzano, A. Cimatti, C. Mattarei, and S. Tonetta. Formal safety assessment via contract-based design. In *Automated Technology for Verification and Analysis*, pages 81–97. Springer, 2014.

[5] C. Bron and J. Kerbosch. Algorithm 457: finding all cliques of an undirected graph. *Communications of the ACM*, 16(9):575–577, 1973.

[6] K. Chatterjee, T. A. Henzinger, J. Otop, and A. Pavlogiannis. Distributed synthesis for ltl fragments. In *Formal Methods in Computer-Aided Design (FMCAD), 2013*, pages 18–25. IEEE, 2013.

[7] D. Cheng, H. Qi, and Z. Li. *Analysis and control of Boolean networks: a semi-tensor product approach*. Springer Science & Business Media, 2010.

[8] L. De Alfaro and T. A. Henzinger. Interface theories for component-based design. In *Embedded Software*, pages 148–165. Springer, 2001.

[9] B. Finkbeiner and S. Schewe. Uniform distributed synthesis. *Logic in Computer Science, Symposium on*, 0:321–330, 2005.

[10] R. A. Hearn and E. D. Demaine. *Games, puzzles, and computation*. CRC Press, 2009.

[11] S. Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22(3):437 – 467, 1969.

[12] P. Madhusudan and P. Thiagarajan. Distributed controller synthesis for local specifications. In F. Orejas, P. Spirakis, and J. van Leeuwen, editors, *Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, pages 396–407. Springer Berlin, 2001.

[13] I. Moir and A. Seabridge. *Aircraft Systems: Mechanical, Electrical, and Avionics Subsystems Integration*. AIAA Education Series, 2001.

[14] P. Nuzzo, H. Xu, N. Ozay, J. B. Finn, A. L. Sangiovanni-Vincentelli, R. M. Murray, A. Donzé, S. Seshia, et al. A contract-based methodology for aircraft electric power system design. *Access, IEEE*, 2:1–25, 2014.

[15] C. M. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.

[16] A. Pneuli and R. Rosner. Distributed reactive systems are hard to synthesize. In *SFCS '90: Proceedings of the 31st Annual Symposium on Foundations of Computer Science*, pages 746–757 vol.2, Washington, DC, USA, 1990. IEEE Computer Society.

[17] Y. E. Sahin and N. Ozay. SAT-based distributed reactive control protocol synthesis for boolean networks. In *IEEE Multi-Conference on Systems and Control (MSC)*, 2016.

[18] Y. E. Sahin and N. Ozay. WiP abstract: Distributed reactive control synthesis for aircraft electric power systems via SAT solving. In *ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*, 2016.

[19] P. Tabuada. *Verification and control of hybrid systems: a symbolic approach*. Springer, 2009.